Refine Search

Search Results -

Terms	Documents
L3 same node	62

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L4

Refine Search

Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, September 07, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=PGPB,USPT,USOC; PLUR=YES; OP=OR			
L4	L3 same node	62	L4
L3	L1 same cache	371	L3
L2	L1 and cache	732	L2
L1	control$4 same ((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	1085	L1

END OF SEARCH HISTORY

Refine Search

Search Results -

Terms	Documents
L4	0

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L5

Refine Search

Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, September 07, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR			
L5	L4	0	L5
DB=PGPB,USPT,USOC; PLUR=YES; OP=OR			
L4	L3 same node	62	L4
L3	L1 same cache	371	L3
L2	L1 and cache	732	L2
L1	control$4 same ((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	1085	L1

END OF SEARCH HISTORY

Refine Search

Search Results -

Terms	Documents
(700/5 |709/213 |709/214 |709/251 |710/305 |710/317 |710/300 |710/62 |710/4 |710/72 |711/141 |711/148 |711/120 |712/14 |712/211).ccls.	5746

Database:
- US Pre-Grant Publication Full-Text Database
- US Patents Full-Text Database
- US OCR Full-Text Database
- EPO Abstracts Database
- JPO Abstracts Database
- Derwent World Patents Index
- IBM Technical Disclosure Bulletins

Search: L6

Refine Search

Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, September 07, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L6	710/305,317,300,62,4,72;711/141,148,120;709/213,214,251;700/5;712/14,211.ccls.	5746	L6
	DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR		
L5	L4	0	L5
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L4	L3 same node	62	L4
L3	L1 same cache	371	L3
L2	L1 and cache	732	L2
L1	control$4 same ((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	1085	L1

END OF SEARCH HISTORY

Refine Search

Search Results -

Terms	Documents
L4 and L6	29

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L7

Refine Search

Recall Text | Clear | Interrupt

Search History

DATE: Tuesday, September 07, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L7	l4 and L6	29	L7
L6	710/305,317,300,62,4,72;711/141,148,120;709/213,214,251;700/5;712/14,211.ccls.	5746	L6
	DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR		
L5	L4	0	L5
	DB=PGPB,USPT,USOC; PLUR=YES; OP=OR		
L4	L3 same node	62	L4
L3	L1 same cache	371	L3
L2	L1 and cache	732	L2
L1	control$4 same ((shared or common) near5 memory) same (multiprocessor or (multi adj1 processor))	1085	L1

END OF SEARCH HISTORY

h e b b cg b e e ch

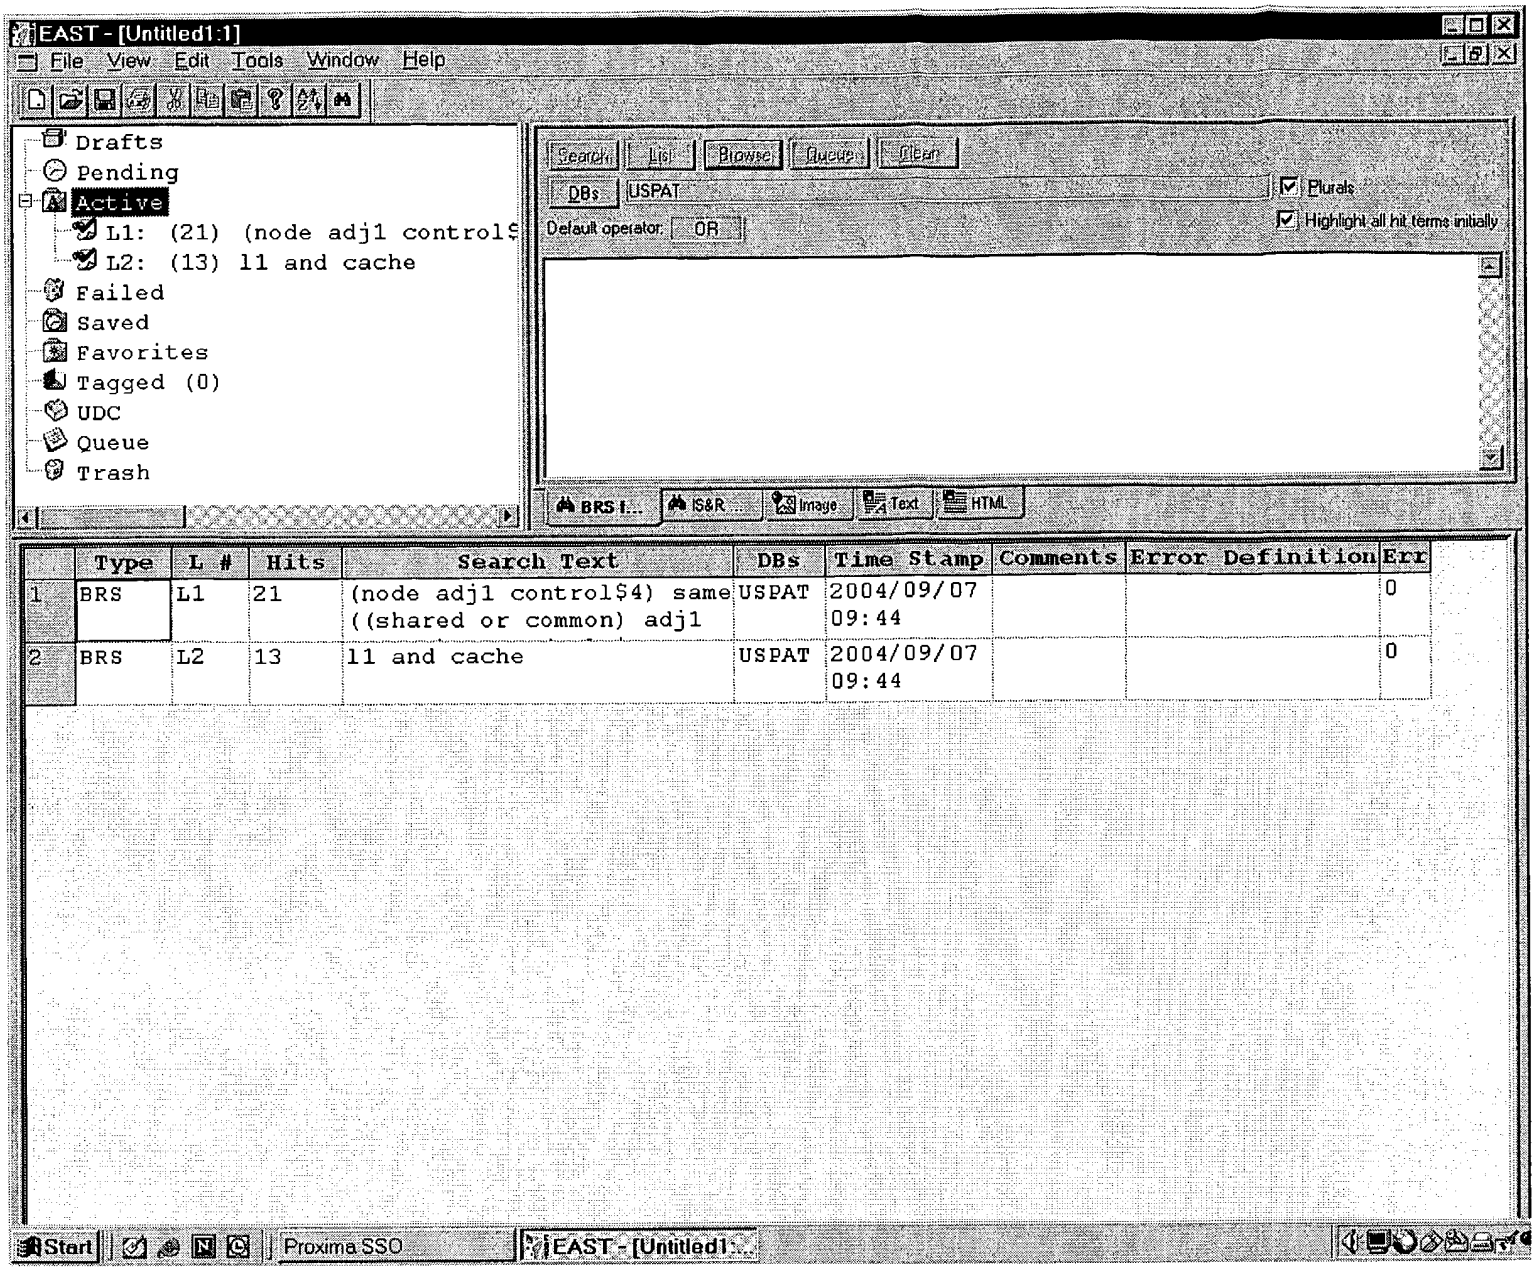

	Type	L #	Hits	Search Text	DBs	Time Stamp	Comments	Error Definition	Err
1	BRS	L1	21	(node adj1 control$4) same ((shared or common) adj1	USPAT	2004/09/07 09:44			0
2	BRS	L2	13	l1 and cache	USPAT	2004/09/07 09:44			0

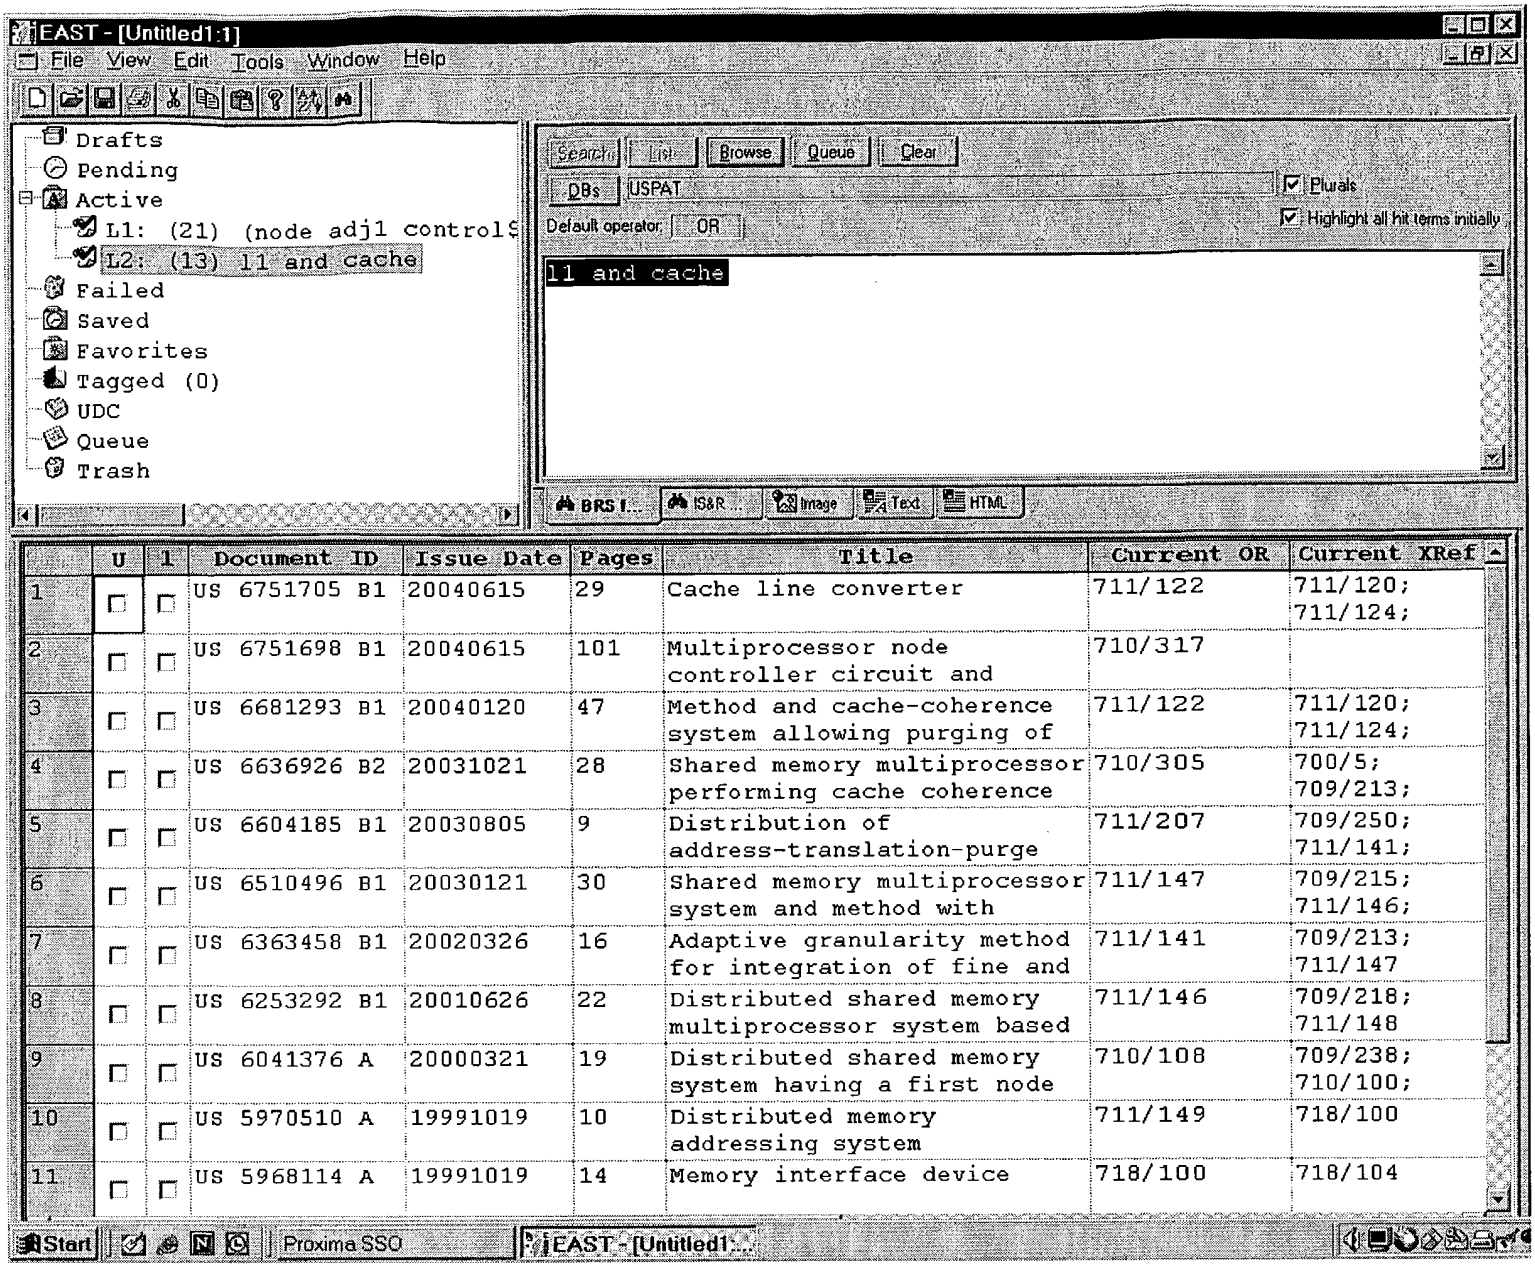

	U	1	Document ID	Issue Date	Pages	Title	Current OR	Current XRef
1	☐	☐	US 6751705 B1	20040615	29	Cache line converter	711/122	711/120; 711/124;
2	☐	☐	US 6751698 B1	20040615	101	Multiprocessor node controller circuit and	710/317	
3	☐	☐	US 6681293 B1	20040120	47	Method and cache-coherence system allowing purging of	711/122	711/120; 711/124;
4	☐	☐	US 6636926 B2	20031021	28	Shared memory multiprocessor performing cache coherence	710/305	700/5; 709/213;
5	☐	☐	US 6604185 B1	20030805	9	Distribution of address-translation-purge	711/207	709/250; 711/141;
6	☐	☐	US 6510496 B1	20030121	30	Shared memory multiprocessor system and method with	711/147	709/215; 711/146;
7	☐	☐	US 6363458 B1	20020326	16	Adaptive granularity method for integration of fine and	711/141	709/213; 711/147
8	☐	☐	US 6253292 B1	20010626	22	Distributed shared memory multiprocessor system based	711/146	709/218; 711/148
9	☐	☐	US 6041376 A	20000321	19	Distributed shared memory system having a first node	710/108	709/238; 710/100;
10	☐	☐	US 5970510 A	19991019	10	Distributed memory addressing system	711/149	718/100
11	☐	☐	US 5968114 A	19991019	14	Memory interface device	718/100	718/104

Your search matched **17** of **1069805** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enter new one in the text box.

(shared or common) and memory and multiprocessc Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

1 **A. novel approach to reduce L2 miss latency in shared-memory multiprocessors**
Acacio, M.E.; Gonzalez, J.; Garcia, J.M.; Duato, J.;
Parallel and Distributed Processing Symposium., Proceedings International, IF 2002, Abstracts and CD-ROM , 15-19 April 2002
Pages:62 - 69

[Abstract] [PDF Full-Text (319 KB)] **IEEE CNF**

---

2 **Flexible use of memory for replication/migration in cache-coherent multiprocessors**
Soundararajan, V.; Heinrich, M.; Verghese, B.; Gharachorloo, K.; Gupta, A.; Hennessy, J.;
Computer Architecture, 1998. Proceedings. The 25th Annual International Symposium on , 27 June-1 July 1998
Pages:342 - 355

[Abstract] [PDF Full-Text (80 KB)] **IEEE CNF**

---

3 **Architectural support for uniprocessor and multiprocessor active memory systems**
Kim, D.; Chaudhuri, M.; Heinrich, M.; Speight, E.;
Computers, IEEE Transactions on , Volume: 53 , Issue: 3 , March 2004
Pages:288 - 307

[Abstract] [PDF Full-Text (1212 KB)] **IEEE JNL**

---

4 **The impact of negative acknowledgments in shared memory scienti applications**
Mainak Chaudhuri; Heinrich, M.;

Parallel and Distributed Systems, IEEE Transactions on , Volume: 15 , Issue: 2 , Feb 2004
Pages:134 - 150

[Abstract] [PDF Full-Text (1831 KB)] IEEE JNL

---

5 **Performance and configuration of hierarchical ring networks for multiprocessors**
Hamacher, V.C.; Hong Jiang;
Parallel Processing, 1997., Proceedings of the 1997 International Conference on , 11-15 Aug. 1997
Pages:257 - 265

[Abstract] [PDF Full-Text (944 KB)] IEEE CNF

---

6 **Design trade-offs in high-throughput coherence controllers**
Nguyen, A.-T.; Torrellas, J.;
Parallel Architectures and Compilation Techniques, 2003. PACT 2003. Proceec 12th International Conference on , 27 Sept.-1 Oct. 2003
Pages:194 - 205

[Abstract] [PDF Full-Text (292 KB)] IEEE CNF

---

7 **Coherent block data transfer in the FLASH multiprocessor**
Heinlein, J.; Bosch, R.P., Jr.; Gharachorloo, K.; Rosenblum, M.; Gupta, A.;
Parallel Processing Symposium, 1997. Proceedings., 11th International , 1-5 , 1997
Pages:18 - 27

[Abstract] [PDF Full-Text (1188 KB)] IEEE CNF

---

8 **SOME-Bus-NOW: a Network of Wrkstations with broadcast**
Katsinis, C.; Hecht, D.;
Network Computing and Applications, 2003. NCA 2003. Second IEEE Internat Symposium on , 16-18 April 2003
Pages:113 - 120

[Abstract] [PDF Full-Text (602 KB)] IEEE CNF

---

9 **Coherence controller architectures for scalable shared-memory multiprocessors**
Michael, M.M.; Nanda, A.K.; Beng-Hong Lim;
Computers, IEEE Transactions on , Volume: 48 , Issue: 2 , Feb. 1999
Pages:245 - 255

[Abstract] [PDF Full-Text (864 KB)] IEEE JNL

---

10 **Toward a cost-effective DSM organization that exploits processor-memory integration**
Torrellas, J.; Liuxi Yang; Nguyen, A.-T.;
High-Performance Computer Architecture, 2000. HPCA-6. Proceedings. Sixth International Symposium on , 8-12 Jan. 2000
Pages:15 - 25

[Abstract] [PDF Full-Text (368 KB)] **IEEE CNF**

11 **PRISM-a design for scalable shared memory**
Kattamuri, E.; Beng-Hong Lim; Pattnaik, P.; Snir, M.;
Innovative Architecture for Future Generation High-Performance Processors a Systems, 1997 , 22-24 Oct. 1997
Pages:29

[Abstract] [PDF Full-Text (76 KB)] **IEEE CNF**

12 **The effect of limited network bandwidth and its utilization by laten hiding techniques in large-scale shared memory systems**
Sunil Kim; Veidenbaum, A.V.;
Parallel Architectures and Compilation Techniques., 1997. Proceedings. 1997 International Conference on , 10-14 Nov. 1997
Pages:40 - 51

[Abstract] [PDF Full-Text (1284 KB)] **IEEE CNF**

13 **Hardware versus software implementation of COMA**
Moga, A.; Gefflaut, A.; Dubois, M.;
Parallel Processing, 1997., Proceedings of the 1997 International Conference on , 11-15 Aug. 1997
Pages:248 - 256

[Abstract] [PDF Full-Text (972 KB)] **IEEE CNF**

14 **Performance evaluation of a WDMA OIDSM multiprocessors**
I-Shyan Hwang;
Parallel and Distributed Systems, 1996. Proceedings., 1996 International Conference on , 3-6 June 1996
Pages:162 - 168

[Abstract] [PDF Full-Text (556 KB)] **IEEE CNF**

15 **TWIN: a parallel scheme for a production system featuring both co and data parallelism**
Yukawa, T.; Ishikawa, T.; Kikuchi, H.; Matsuzawa, K.;
Artificial Intelligence for Applications, 1991. Proceedings., Seventh IEEE Conf on , Volume: i , 24-28 Feb. 1991
Pages:64 - 70

[Abstract] [PDF Full-Text (500 KB)] **IEEE CNF**

1 2 Next

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account | New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online Publications | Help | FAQ | Terms | Back to Top

Architectural support for uniprocessor and multiprocessor active memory systems

Kim, D. Chaudhuri, M. Heinrich, M. Speight, E.
Comput. Syst. Lab., Cornell Univ., Ithaca, NY, USA

This paper appears in: **Computers, IEEE Transactions on**

Publication Date: March 2004
On page(s): 288 - 307
Volume: 53 , Issue: 3
ISSN: 0018-9340
Inspec Accession Number: 8004231

Abstract:
We introduce an architectural approach to improve **memory** system performance in both uniprocessor and **multiprocessor** systems. The architectural innovation is a flexible active **memory controller** backed by specialized **cache** coherence protocols that permit the transparent use of address remapping techniques. The resulting system shows significant performance improvement across a spectrum of machine configurations, from uniprocessors through single-**node multiprocessors** (SMPs) to distributed **shared memory** clusters (DSMs). Address remapping techniques exploit the data access patterns of applications to enhance their **cache** performance. However, they create coherence problems since the processor is allowed to refer to the same data via more

h eee e eee g e ch ch b c be be be c e eee

Print Format

than one address. While most active **memory** implementations require **cache** flushes, we present a new approach to solve the coherence problem. We leverage and extend the **cache** coherence protocol so that our techniques work transparently to efficiently support uniprocessor, SMP and DSM active **memory** systems. We detail the coherence protocol extensions to support our active **memory** techniques and present simulation results that show uniprocessor speedup from 1.3 to 7.6 on a range of applications and microbenchmarks. We also show remarkable performance improvement on small to medium-scale SMP and DSM **multiprocessors**, allowing some parallel applications to continue to scale long after their performance levels off on normal systems.

Index Terms:

cache storage distributed **shared memory** systems **memory** architecture protocols DSM SMP active **memory controller** address remapping technique architectural support **cache** coherence protocol data access pattern distributed **shared memory** cluster machine configuration spectrum **memory** system performance **multiprocessor** system single-**node multiprocessor** uniprocessor system

Documents that cite this document

There are no citing documents available in IEEE Xplore at this time.

Search Results [PDF FULL-TEXT 1212 KB] PREV NEXT DOWNLOAD CITATION

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account | New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online Publications | Help | FAQ | Terms | Back to Top

h eee e eee g e ch ch b c be be be c e eee

A. novel approach to reduce L2 miss latency in shared-memory multiprocessors

Acacio, M.E. Gonzalez, J. Garcia, J.M. Duato, J.
Dpto. Ing. y Tecnologia de Computadores, Murcia Univ., Spain;

This paper appears in: **Parallel and Distributed Processing Symposium., Proceedings International, IPDPS 2002, Abstracts and CD-ROM**

Meeting Date: 04/15/2002 - 04/19/2002
Publication Date: 15-19 April 2002
Location: Ft. Lauderdale, FL USA
On page(s): 62 - 69
Reference Cited: 16
Number of Pages: CD-ROM
Inspec Accession Number: 7342351

Abstract:
Recent technology improvements allow **multiprocessor** designers to put some key components inside the processor chip, such as the **memory controller**, the coherence hardware and the network interface/router. In this work we exploit such integration scale, presenting a novel **node** architecture aimed at reducing the long L2 miss latencies and the **memory** overhead of using directories that characterize cc-NUMA machines and limit their scalability. Our proposal replaces the traditional directory with a novel three-

Print Format

level directory architecture and adds a small **shared** data **cache** to each of the **nodes** of a **multiprocessor** system. Due to their small size, the first-level directory and the **shared** data **cache** are integrated into the processor chip in every **node.** A taxonomy of the L2 misses, according to the actions performed by the directory to satisfy them is also presented. Using execution-driven simulations, we show significant L2 miss latency reductions (more than 60% in some cases). These important improvements translate into reductions of more than 30% in the application execution time in some cases

Index Terms:

cache storage parallel architectures performance evaluation **shared memory** systems L2 miss latency reduction cc-NUMA machines coherence hardware execution-driven simulations **memory controller** **memory** overhead network interface **node** architecture scalability **shared** data **cache** **shared-memory multiprocessors** three-level directory architecture

Documents that cite this document

There are no citing documents available in IEEE Xplore at this time.

Search Results [PDF FULL-TEXT 319 KB] NEXT DOWNLOAD CITATION

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account | New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online Publications | Help | FAQ| Terms | Back to Top

h eee e eee g e ch ch b c be be be c e e

L7: Entry 5 of 29 File: PGPB Feb 19, 2004

PGPUB-DOCUMENT-NUMBER: 20040034747
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20040034747 A1

TITLE: Scalable cache coherent distributed shared memory processing system

PUBLICATION-DATE: February 19, 2004

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Rowlands, Joseph B.	Santa Clara	CA	US	
Gulati, Manu	Santa Clara	CA	US	

APPL-NO: 10/ 356321 [PALM]
DATE FILED: January 31, 2003

RELATED-US-APPL-DATA:
Application is a non-provisional-of-provisional application 60/380740, filed May 15, 2002,
Application is a non-provisional-of-provisional application 60/419033, filed October 16, 2002,

INT-CL: [07] G06 F 12/08

US-CL-PUBLISHED: 711/148; 711/119, 711/144, 711/145
US-CL-CURRENT: 711/148; 711/119, 711/144, 711/145

REPRESENTATIVE-FIGURES: 14

ABSTRACT:

A packetized I/O link such as the HyperTransport protocol is adapted to transport memory coherency transactions over the link to support cache coherency in distributed shared memory systems. The I/O link protocol is adapted to include additional virtual channels that can carry command packets for coherency transactions over the link in a format that is acceptable to the I/O protocol. The coherency transactions support cache coherency between processing nodes interconnected by the link. Each processing node may include processing resources that themselves share memory, such as symmetrical multiprocessor configuration. In this case, coherency will have to be maintained both at the intranode level as well as the internode level. A remote line directory is maintained by each processing node so that it can track the state and location of all of the lines from its local memory that have been provided to other remote nodes. A node controller initiates transactions over the link in response to local transactions initiated within itself, and initiates transactions over the link based on local transactions initiated within itself. Flow control is provided for each of the coherency virtual channels either by software through credits or through a buffer free command packet

that is sent to a source node by a target node indicating the availability of virtual channel buffering for that channel.

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims priority under 35 U.S.C. 119(e) to the following applications, each of which is incorporated herein for all purposes:

[0002] (1) provisional patent application entitled SYSTEM ON A CHIP FOR NETWORKING, having an application No. 60/380,740, and a filing date of May 15, 2002; and

[0003] (2) provisional patent application having the same title as above, having an application No. 60/419,033, and a filing date of Oct. 16, 2002.

Previous Doc Next Doc Go to Doc#

L7: Entry 14 of 29 File: PGPB Jun 28, 2001

PGPUB-DOCUMENT-NUMBER: 20010005873
PGPUB-FILING-TYPE: new-utility
DOCUMENT-IDENTIFIER: US 20010005873 A1

TITLE: Shared memory multiprocessor performing cache coherence control and node controller therefor

PUBLICATION-DATE: June 28, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Yasuda, Yoshiko	Tokorozawa		JP	
Hamanaka, Naoki	Tokyo		JP	
Shonai, Toru	Hachioji		JP	
Akashi, Hideya	Kunitachi		JP	
Tsushima, Yuji	Kokubunji		JP	
Uehara, Keitaro	Kokubunji		JP	

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	COUNTRY	TYPE CODE
Hitachi, Ltd.				03

APPL-NO: 09/ 740816 [PALM]
DATE FILED: December 21, 2000

FOREIGN-APPL-PRIORITY-DATA:

COUNTRY	APPL-NO	DOC-ID	APPL-DATE
JP	11-366235	1999JP-11-366235	December 24, 1999

INT-CL: [07] G06 F 13/00, G06 F 13/38

US-CL-PUBLISHED: 710/129
US-CL-CURRENT: 710/305

REPRESENTATIVE-FIGURES: 1

ABSTRACT:

Each node includes a node controller for decoding the control information and the address information for the access request issued by a processor or an I/O device, generating, based on the result of decoding, the cache coherence control information indicating whether the cache coherence control is required or not, the node information and the unit information for the transfer destination, and adding these information to the access request. An intra-node connection circuit for

connecting the units in the node controller holds the cache coherence control information, the node information and the unit information added to the access request. When the cache coherence control information indicates that the cache coherence control is not required and the node information indicates the local node, then the intra-node connection circuit transfers the access request not to the inter-node connection circuit interconnecting the nodes but directly to the unit designated by the unit information.

Previous Doc Next Doc Go to Doc#

L7: Entry 21 of 29 File: USPT Oct 21, 2003

US-PAT-NO: 6636926
DOCUMENT-IDENTIFIER: US 6636926 B2

TITLE: Shared memory multiprocessor performing cache coherence control and node controller therefor

DATE-ISSUED: October 21, 2003

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Yasuda; Yoshiko	Tokorozawa			JP
Hamanaka; Naoki	Tokyo			JP
Shonai; Toru	Hachioji			JP
Akashi; Hideya	Kunitachi			JP
Tsushima; Yuji	Kokubunji			JP
Uehara; Keitaro	Kokubunji			JP

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Hitachi, Ltd.	Tokyo			JP	03

APPL-NO: 09/ 740816 [PALM]
DATE FILED: December 21, 2000

FOREIGN-APPL-PRIORITY-DATA:

COUNTRY	APPL-NO	APPL-DATE
JP	11-366235	December 24, 1999

INT-CL: [07] G06 F 13/00, G06 F 15/167

US-CL-ISSUED: 710/305; 710/317, 711/141, 709/213, 700/5
US-CL-CURRENT: 710/305; 700/5, 709/213, 710/317, 711/141

FIELD-OF-SEARCH: 710/305, 710/317, 710/300, 710/62, 710/4, 710/72, 711/141, 711/148, 711/120, 709/213, 709/214, 709/251, 700/5, 712/14, 712/211

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL

h e b b g eeef c e e ge

☐	4747043	May 1988	Rodman
☐	6011791	January 2000	Okada et al.
☐	6092173	July 2000	Sasaki et al.
☐	6378029	April 2002	Venkitakrishnan et al.
☐	6466825	October 2002	Wang et al.

OTHER PUBLICATIONS

"RISC System/6000SMP System," 1995 Comcon95 Proceedings, pp. 102-109.
"Starfire: Extending the SMP Envelope," 1998 Micro Jan./Feb. pp. 39-49.

ART-UNIT: 2181

PRIMARY-EXAMINER: Ray; Gopal C.

ABSTRACT:

Each node includes a node controller for decoding the control information and the address information for the access request issued by a processor or an I/O device, generating, based on the result of decoding, the cache coherence control information indicating whether the cache coherence control is required or not, the node information and the unit information for the transfer destination, and adding these information to the access request. An intra-node connection circuit for connecting the units in the node controller holds the cache coherence control information, the node information and the unit information added to the access request. When the cache coherence control information indicates that the cache coherence control is not required and the node information indicates the local node, then the intra-node connection circuit transfers the access request not to the inter-node connection circuit interconnecting the nodes but directly to the unit designated by the unit information.

19 Claims, 16 Drawing figures

Previous Doc Next Doc Go to Doc#

First Hit Fwd Refs Previous Doc Next Doc Go to Doc#

Generate Collection Print

L7: Entry 23 of 29 File: USPT Apr 8, 2003

US-PAT-NO: 6546471
DOCUMENT-IDENTIFIER: US 6546471 B1

TITLE: Shared memory multiprocessor performing cache coherency

DATE-ISSUED: April 8, 2003

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Tarui; Toshiaki	Sagamihara			JP
Okazawa; Koichi	Ebina			JP
Okada; Yasuyuki	Yamato			JP
Shonai; Toru	Kodaira			JP
Okochi; Toshio	Kokubunji			JP
Akashi; Hideya	Hachiouji			JP

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Hitachi, Ltd.	Tokyo			JP	03

APPL-NO: 09/ 506810 [PALM]
DATE FILED: February 18, 2000

PARENT-CASE:
This is a continuation application of U.S. Ser. No. 09/030,957, filed Feb. 26, 1998, now U.S. Pat. No. 6,088,770.

FOREIGN-APPL-PRIORITY-DATA:

COUNTRY	APPL-NO	APPL-DATE
JP	9-059914	February 27, 1997

INT-CL: [07] G06 F 13/14

US-CL-ISSUED: 711/148; 711/141, 711/147, 711/149, 711/169
US-CL-CURRENT: 711/148; 711/141, 711/147, 711/149, 711/169

FIELD-OF-SEARCH: 711/100, 711/113, 711/118, 711/119, 711/133-135, 711/141-149, 711/169, 711/206, 709/200, 709/201, 709/213, 709/218, 364/131

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

h e b b g e e e f c e e ge

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4985825	January 1991	Webb, Jr. et al.	711/169
☐	5117350	May 1992	Parrish et al.	711/1
☐	5710907	January 1998	Hagersten et al.	711/148
☐	5829034	October 1998	Hagersten et al.	711/141
☐	5890189	March 1999	Nozue et al.	711/100
☐	5956754	September 1999	Kimmel	711/206
☐	6456628	September 2002	Greim et al.	370/466

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	US-CL
08-320828	March 1996	JP	
09-022381	January 1997	JP	

OTHER PUBLICATIONS

"Evolved System Architecture", Sun World, Jan. 1996, pp. 29-32.
"The Stanford Flash Multiprocessor", (The 21st Annual International Symposium on Computer Architecture, Apr. 18-21, 1994, Chicago, Illinois, pp. 302-313.
"Hive: Fault Containment for Shared-Memory Multiprocessors" (15th ACM Symposium on Operating Systems Principles, Dec. 3-6, 1995, Copper Mountain Resort, Colorado, pp. 12-25.
"DDM-A Cache-Only Memory Architecture", Computer, Sep. 1992, vol. 25, pp. 44-54.
"Cache coherent shared memory hypercube", Parallel and Distributed Processing, 1992. Proceedings of the 4th IEEE Symposium, Dec. 1-4 1992, pp. 515-520.
"Software Cachingon cache-coherent multiprocessors", Parallel and Distributed Processing, 1992. Proceedings of the 4th IEEE Symposium, Dec. 1-4 1992, pp. 521-526.

ART-UNIT: 2187

PRIMARY-EXAMINER: Nguyen; T. V.

ATTY-AGENT-FIRM: Mattingly, Stanger & Malur, P.C.

ABSTRACT:

A shared memory multiprocessor (SMP) has efficient access to a main memory included in a particular node and a management of partitions that include the nodes. In correspondence with each page of main memory included in a node, a bit stored in a register indicates if the page has been accessed from any other node. In a case where the bit is "0", a cache coherent command to be sent to the other nodes is not transmitted. The bit is reset by software at the time of initialization and memory allocation, and it is set by hardware when the page of the main memory is accessed from any other node. In a case where the interior of an SMP is divided into partitions, the main memory of each node is divided into local and shared areas, for which respectively separate addresses can be designated. In each node, the configuration information items of the shared area and the local area are stored in registers. The command of access to the shared area is multicast to all of the nodes, whereas the command is multicast only to the nodes within the corresponding

partition when the local area is accessed.

4 Claims, 20 Drawing figures

Previous Doc Next Doc Go to Doc#

(12) **United States Patent**
Deneroff et al.

(10) **Patent No.: US 6,751,698 B1**
(45) **Date of Patent: Jun. 15, 2004**

(54) **MULTIPROCESSOR NODE CONTROLLER CIRCUIT AND METHOD**

(75) Inventors: **Martin M. Deneroff**, Palo Alto, CA (US); **Givargis G. Kaldani**, San Jose, CA (US); **Yuval Koren**, San Francisco, CA (US); **David Edward McCracken**, San Francisco, CA (US); **Swami Venkataraman**, San Jose, CA (US)

(73) Assignee: **Silicon Graphics, Inc.**, Mountain View, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/407,428**

(22) Filed: **Sep. 29, 1999**

(51) **Int. Cl.**[7] **G06F 13/00**
(52) **U.S. Cl.** **710/317**
(58) **Field of Search** 710/100, 305, 710/306, 311, 316, 317; 709/238, 218, 249; 712/12; 370/351, 419

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,330,858	A	5/1982	Choquet	370/111
4,630,259	A	12/1986	Larson et al.	370/60

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

EP	0501524	9/1992	G06F/15/16
EP	0570729	11/1993	G06F/15/16
WO	96/37834	11/1996	G06F/9/46
WO	99/26429	5/1999	H04Q/3/00

OTHER PUBLICATIONS

International Search Report for International Application No. PCT/IE95/00047; Date of Completion—Dec. 22, 1995; Authorized Officer—R. Salm.

"Deadlock–Free Routing Schemes on Multistage Interconnection Networks", *IBM Technical Disclosure Bulletin*, 35, 232–233, (Dec., 1992).

"IEEE Standard for Scalable Coherent Interface (SCI)", *IEEE Std 1596–1992*, Table of Contents, (Mar., 1992).

(List continued on next page.)

Primary Examiner—Xuan M. Thai
(74) *Attorney, Agent, or Firm*—Schwegman, Lundberg, Woessner & Kluth, P.A.

(57) **ABSTRACT**

Improved method and apparatus for parallel processing. One embodiment provides a multiprocessor computer system that includes a first and second node controller, a number of processors being connected to each node controller, a memory connected to each controller, a first input/output system connected to the first node controller, and a communications network connected between the node controllers. The first node controller includes: a crossbar unit to which are connected a memory port, an input/output port, a network port, and a plurality of independent processor ports. A first and a second processor port connected between the crossbar unit and a first subset and a second subset, respectively, of the processors. In some embodiments of the system, the first node controller is fabricated onto a single integrated-circuit chip. Optionally, the memory is packaged on plugable memory/directory cards wherein each card includes a plurality of memory chips including a first subset dedicated to holding memory data and a second subset dedicated to holding directory data. Further, the memory port includes a memory data port including a memory data bus and a memory address bus coupled to the first subset of memory chips, and a directory data port including a directory data bus and a directory address bus coupled to the second subset of memory chips. In some such embodiments, the ratio of (memory data space) to (directory data space) on each card is set to a value that is based on a size of the multiprocessor computer system.

31 Claims, 70 Drawing Sheets

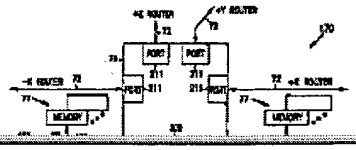

(12) **United States Patent**
Fromm

(10) **Patent No.: US 6,604,185 B1**
(45) **Date of Patent: Aug. 5, 2003**

(54) **DISTRIBUTION OF ADDRESS-TRANSLATION-PURGE REQUESTS TO MULTIPLE PROCESSORS**

(75) Inventor: **Eric C. Fromm**, Eau Claire, WI (US)

(73) Assignee: **Silicon Graphics, Inc.**, Mountain View, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/619,851**

(22) Filed: **Jul. 20, 2000**

(51) **Int. Cl.**[7] **G06F 12/12**

(52) **U.S. Cl.** **711/207**; 711/141; 711/166; 709/250

(58) **Field of Search** 709/1, 213, 218, 709/232, 250; 711/141, 207, 206, 202, 166; 713/2; 710/53, 54

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,779,188	A	* 10/1988	Gum et al.	709/1
5,437,017	A	* 7/1995	Moore et al.	709/213
5,506,953	A	4/1996	Dao	
5,574,878	A	* 11/1996	Onodera et al.	711/207
5,603,056	A	2/1997	Totani	
5,617,537	A	4/1997	Yamada	
5,644,575	A	* 7/1997	McDaniel	370/416
5,682,512	A	10/1997	Tetrick	
577,429	A	7/1998	Sukegawa et al	711/129
5,784,394	A	* 7/1998	Alvarez et al.	714/799
5,784,706	A	7/1998	Oberlin et al	711/202
5,906,001	A	* 5/1999	Wu et al.	711/141
6,128,282	A	* 10/2000	Liebetreu et al.	370/235
6,339,812	B1	* 1/2002	McCracken et al.	711/141
6,374,331	B1	* 4/2002	Janakiraman et al.	711/141
6,412,056	B1	6/2002	Ghanchorloo et al	711/202

OTHER PUBLICATIONS

Gjessing, et al., "Performance of the RamLink Memory Architecture", *Proceedings HICSS'94*, (1994), 154–162.

Gjessing , Stein., et al., "RamLink: A High–Bandwidth-Point–to–Point Memory Architecture", *Proceeding Comp-Con*, (1992),328/331.

IEEE STD, "IEEE Standard for High–Bandwidth Memory Interface Based on Scalable Coherent Interface(SCI) Signaling Technology (RAMLink)", *IEEE Std 1596.4–1996*, (1996), 1–91.

* cited by examiner

Primary Examiner—B. James Peikari
(74) *Attorney, Agent, or Firm*—Schwegman, Lundberg, Woessner & Kluth, P.A.

(57) **ABSTRACT**

A method and apparatus for deallocating memory in a multi-processor, shared memory system. In one aspect, a node in the system has a node controller that contains sequencing logic. The sequencing logic receives a command across a network. The sequencing logic translates the received command into a Purge Translation Cache (PTC) instruction and sends the PTC instruction across a bus to a processor. The processor contains bus control logic that receives the PTC instruction and purges a virtual address specified in the PTC instruction from the processor's translation lookaside buffer. By purging the virtual address, the memory is deallocated.

16 Claims, 4 Drawing Sheets

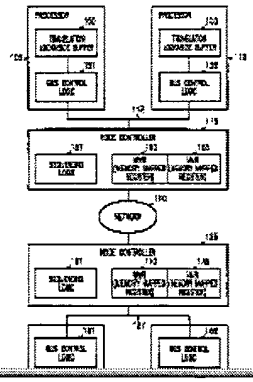

(12) **United States Patent**
Jhang et al.

(10) **Patent No.: US 6,253,292 B1**
(45) **Date of Patent: Jun. 26, 2001**

(54) **DISTRIBUTED SHARED MEMORY MULTIPROCESSOR SYSTEM BASED ON A UNIDIRECTIONAL RING BUS USING A SNOOPING SCHEME**

(76) Inventors: **Seong Tae Jhang**, 103-1611, Wooseong Apt., San 72, Jeongrung 1-Dong, Seongbuk-Gu, Seoul (KR), 136-101; **Chu Shik Jhon**, 125-601, Keonyoung Apt., Jangan Town, Pundang-Dong, Pundang-Gu, Seongnam-Shi, Kyoungki-Do (KR), 463-030; **Hyung Ho Kim**, 504-1108, Jukong Apt., Byoulyang-Dong, Guacheon-Shi, Kyoungki-Do (KR), 427-040

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/120,850**

(22) Filed: **Jul. 22, 1998**

(30) **Foreign Application Priority Data**

Aug. 22, 1997 (KR) 97-40083
Apr. 23, 1998 (KR) 98-14513

(51) **Int. Cl.**[7] **G06F 12/00**
(52) **U.S. Cl.** **711/146**; 711/148; 709/218
(58) **Field of Search** 711/141, 145, 711/147, 148; 710/128; 709/216, 217, 218, 251

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,566,177 * 10/1996 Bhandari et al. 370/85.5
5,588,131 * 12/1996 Borrill 395/473
5,781,757 * 7/1998 Deshpande 395/473

* cited by examiner

Primary Examiner—Kevin Verbrugge
(74) *Attorney, Agent, or Firm*—Anderson Kill & Olick, PC

(57) **ABSTRACT**

A distributed shared memory multiprocessor system based on a unidirectional ring bus using a snooping scheme comprises a group of processor nodes and a ring bus. The processor nodes are arranged in the form of a ring and one of the processor nodes generates a request signal for a data block, the remaining processor nodes snoop their own internal parts, and one of the remaining processor nodes provides the data block. The ring bus is used for connecting the processor nodes in the form of the ring and providing a path through which the request signal is broadcast to each of the remaining processor nodes and the data block is unicast to the processor node which has generated the request signal for the data block.

13 Claims, 13 Drawing Sheets

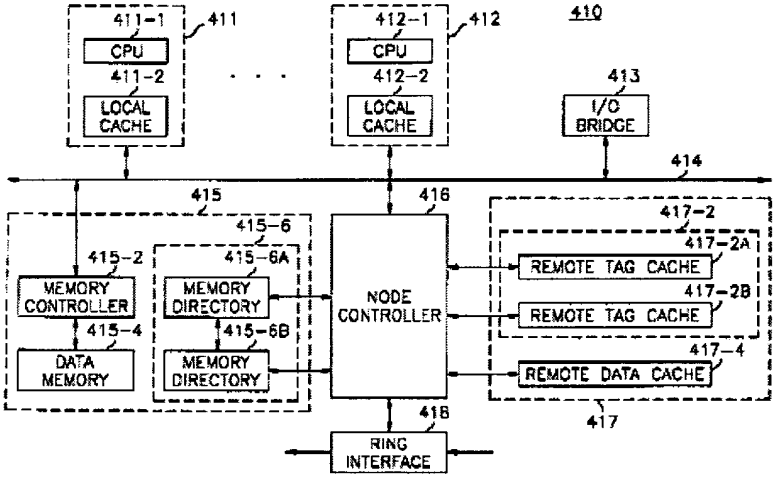

[54] **DISTRIBUTED MEMORY ADDRESSING SYSTEM**

[75] Inventors: **Richard A. Sher**, Huntington; **Jerry Rogers**, Seaford; **Mark J. Wentka**, East Northport, all of N.Y.

[73] Assignee: **Northrop Grumman Corporation**, Los Angeles, Calif.

[*] Notice: This patent issued on a continued prosecution application filed under 37 CFR 1.53(d), and is subject to the twenty year patent term provisions of 35 U.S.C. 154(a)(2).

[21] Appl. No.: 08/629,839

[22] Filed: Apr. 10, 1996

[51] Int. Cl.[6] G06F 9/45

[52] U.S. Cl. 711/149; 709/100

[58] Field of Search 709/102, 100, 709/800.11, 106; 711/149, 203, 202

[56] References Cited

U.S. PATENT DOCUMENTS

4,803,685	2/1989	Rypinski	340/825.05
4,972,314	11/1990	Getzinger	711/149
5,325,526	6/1994	Cameron et al.	395/672
5,384,773	1/1995	Olnowich et al.	365/189.01
5,535,373	7/1996	Olnowich	395/500
5,696,922	12/1997	Froman	711/5

Primary Examiner—Majid A. Banankhah

Attorney, Agent, or Firm—Terry J. Anderson; Karl J. Hoch, Jr.

[57] **ABSTRACT**

A distributed memory addressing system has a plurality of separate processing elements. Each processing element has at least one CPU. A shared memory is utilized to store data to be used by the separate processing elements, as required. A high bandwidth interface interconnects processing elements and the shared memory. The high bandwidth interface is configured so as to provide non-blocking access to the shared memory for each of the processing elements.

25 Claims, 3 Drawing Sheets

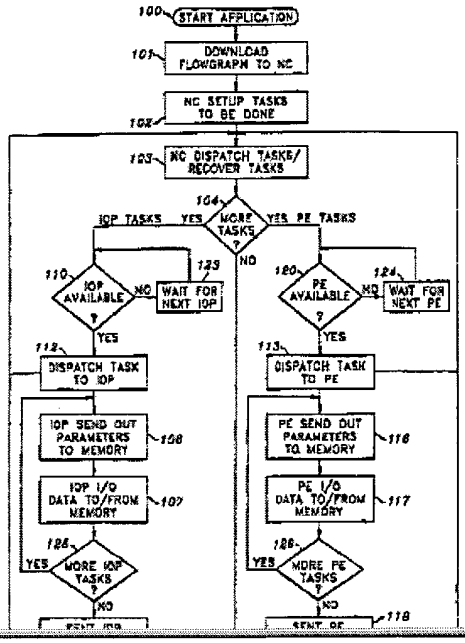

[54] **MEMORY INTERFACE DEVICE**

[75] Inventors: **Mark J. Wentka**, E. Northport; **Richard A. Sher**, Huntington, both of N.Y.

[73] Assignee: **Northrop Grumman Corporation**, Los Angeles, Calif.

[21] Appl. No.: **08/822,746**

[22] Filed: **Mar. 24, 1997**

Related U.S. Application Data

[63] Continuation of application No. 08/629,839, Apr. 10, 1996.

[51] **Int. Cl.**[6] **G06F 9/00**

[52] **U.S. Cl.** **709/100**; 709/104

[58] **Field of Search** 709/100, 102, 709/104, 105, 106; 711/100, 170; 712/38, 16

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,177,511	12/1979	Taddel	364/200
4,454,575	6/1984	Bushaw et al.	364/200
4,476,524	10/1984	Brown et al.	710/126
4,787,041	11/1988	Yount	364/424
4,821,185	4/1989	Esposito	364/200
4,847,757	7/1989	Smith	710/114
4,901,230	2/1990	Chen et al.	364/200
5,237,673	8/1993	Orbits et al.	395/425
5,247,630	9/1993	Parks et al.	395/400
5,388,841	2/1995	San et al.	463/44
5,390,336	2/1995	Hillis	395/800
5,410,654	4/1995	Foster et al.	395/275
5,458,666	8/1995	Craft et al.	395/842
5,466,862	8/1995	Ohkami	711/100
5,475,858	12/1995	Gupta et al.	395/800
5,485,590	1/1996	Hyatt et al.	395/442
5,487,146	1/1996	Guttag et al.	345/516
5,513,374	4/1996	Baji	395/846
5,524,265	6/1996	Balmer et al.	712/38
5,560,030	9/1996	Guttag et al.	712/16
5,850,632	12/1998	Robertson	711/170

Primary Examiner—Majid A. Banankhah
Attorney, Agent, or Firm—Terry J. Anderson; Karl J. Hoch, Jr.

[57] **ABSTRACT**

A memory interface device for facilitating electrical communication between distributed memory and a plurality of processors has a memory interface circuit configured to interface the memory interface device to at least one random access memory, an address generator circuit configured to generate addresses for data stored within the random access memories and a processor interface circuit configured to interface the memory interface device to a plurality of processors. Interfacing the memory interface device to both the random access memories and the plurality of processors facilitates simultaneous non-interruptible access by all of the processors to data stored in the random access memories.

17 Claims, 4 Drawing Sheets